Fairness in Decentralized Strategic Deconfliction in UTM

Antony Evans, PhD;[1] Maxim Egorov[2] and Steven Munn[3]
Airbus UTM, Sunnyvale, CA, 94086

Decentralized strategic deconfliction is a critical function in Unmanned Aircraft System (UAS) Traffic Management (UTM), that serves as the enabler of safe operations for cooperative traffic. In this paper we explore the implications of different approaches to strategic decentralized trajectory deconfliction in UTM on fairness. First-come, first-served (FCFS) allocation of resources has long-standing acceptance in the aviation community and is applied both tactically and strategically in air traffic management. However, it is unclear what the fairness implications are for a FCFS allocation of access to airspace in a decentralized UTM architecture. In this paper, simulation is used to explore how a FCFS approach to strategic deconfliction in UTM – based on when operators file their flight plans – performs in terms of fairness. Fairness is quantified by comparing average ground delay across operators and by calculating a normalized fairness metric that accounts for operator cost of delay. Two scenario types are simulated: (1) two package delivery operators serving a common region from separate warehouses; and (2) two air taxi operators serving the same network of 7 vertiports. Results are generated for a range of traffic densities and differences in file-ahead time between operators. Results indicate that, for a decentralized FCFS approach to strategic deconfliction based on when operators file their flight plans, there may be a significant imbalance in delays between operators based on how far in advance they are able to file ahead, and on traffic demand levels. To ensure fairness at envisioned traffic densities it may therefore be necessary to constrain how early flight plan requests can be prioritized over later requests – similar in concept to a freeze horizon in TBFM.

I. Introduction

APPLICATIONS for unmanned aircraft systems (UASs) are diverse, and forecasts suggest that the number of UAS in operation will increase dramatically over the next 20 years.[1,2,3] Much of this growth is expected in urban environments, with package delivery operations[2] and vertical take-off and landing (VTOL) passenger air taxis[3] under development by a number of companies. Such growth in Urban Air Mobility (UAM) will require new approaches to managing air traffic beyond what is currently available through traditional air traffic control (ATC).

A number of proposals exist for modernizing the traffic management process for UAS under a single framework[4,5] – referred to in the United States as UAS Traffic Management (UTM).[4] A key property of the NASA model,[4] also adopted by the U.S. Federal Aviation Administration (FAA),[6] is a decentralized, or federated, architecture, in which UAS Service Suppliers (USSs) provide various services to UAS operators, including managing traffic. The U-Space model, developed by the Single European Sky ATM Research Joint Undertaking (SESAR JU), also provides for a partially decentralized architecture.[5] While the NASA UTM model is focused on small UAS operations, the service-oriented architecture may also be a foundational piece of traffic management for UAM operations, and ultimately the future of traditional Air Traffic Management (ATM).[7]

---

[1] Traffic Management System Architect, Airbus UTM, Senior member.
[2] Research Scientist, Airbus UTM.
[3] Simulation Engineer, Airbus UTM.

UTM operations are rapidly advancing in complexity around the world, as technological capabilities move from simulation to real-world operations[8] and as standards begin to take shape.[9] Numerous demonstrations and trial projects around the world have begun to validate the capabilities of UTM providers to manage multiple aircraft.[10] In the United States, several UAS operators have received their Part 135 air carrier certificates, a key regulatory milestone that enables payload-carrying commercial operations.[11] Prior to these advances, many operations took place in a relative vacuum, or under exemptions issued by the competent authorities. Even in operational trials, conflicts and complexity had been carefully controlled to test core capabilities.[10] Now that traffic numbers are increasing and complex operations are becoming a reality, a unique opportunity exists to refine UTM architectures to ensure that the system designs, implementations and operating rules remain fair for all users, rather than benefiting first-movers and the largest operators or USSs.

Fairness and equity are a key principle in ICAO's UTM Framework,[12] which states that "access to the airspace should remain equitable." This view on fairness is also supported by the FAA UTM Conops,[6] where UTM must "maintain fair and equitable access to airspace," and the European U-Space Conops,[5] where UTM must "guarantee equitable and fair access to airspace for all users."

The current UTM environment lacks a clear and agreed upon set of mechanisms, protocols, or services for resolving competition or conflicts for resources, such as airspace, vertiports, or air traffic services. This has not been an operational problem yet, because the number of operations has not warranted it. But the inability to resolve conflicting requests for resources will soon be an impediment to integration of new entrants. Of particular concern is establishing a UTM environment in which flight operators and other stakeholders view their access to airspace and related resources as fair.

While most industry participants are outwardly collaborative and purport to be "good actors," history shows that, faced with the competitive pressure of open markets, participants will seek ways to gain an unfair advantage, in some cases unfairly. This can have severe consequences for the ultimate goal of providing access to airspace and enabling a wide variety of future missions.

One key area in which fair allocation of resources is of particular concern is the deconfliction of flight operations. The detection and resolution of potential conflicts between vehicles is central to any traffic management function. However, in a federated architecture, in which deconfliction services are distributed, conflict resolution actions need to be negotiated between USSs instead of prescribed by a centralized authority, as in traditional ATM. This is the model applied in the NASA UTM concept, where conflicts identified by strategic deconfliction services can be resolved by peer-to-peer negotiation.[13] While there is a requirement that known conflicts be resolved prior to departure, the rules governing the negotiation process are currently under development within industry standards groups, and have not yet been validated. Ultimately, the rules that are set to govern deconfliction – which are likely to be different for strategic deconfliction (pre-departure), tactical deconfliction (while airborne), and detect-and-avoid systems – could have important implications for the fair allocation of resources.

In this paper we explore existing approaches to allocate resources in traditional ATM and other domains, and existing work in this area in UTM (Section II). This is followed by a description of objectives to quantify fairness implications of a first-come, first-served allocation of resources in a decentralized UTM architecture in Section III. The experimental approach is described in Section IV, followed by results in Section V and conclusions for future work in Section VI.

II. Background

Fairness in the allocation of resources has been studied extensively by economists and mathematicians across a diverse range of domains.[14,15] Lessons learned from a diverse set of domains are presented below, including ATM, UTM and non-aviation domains.

Air Traffic Management

In controlled airspace, the air navigation service provider (ANSP) has strict control over the use of public resources, such as airspace and runways. This means that the processes for allocation are centrally controlled by the ANSP, airport authority or civil aviation authority (CAA). This standard is in contrast to proposals for a more federated architecture in UTM. In traditional ATM, control is distributed hierarchically, typically including a flow control or command center level, with regional ATM services below that. In the U.S., these are radar air traffic control facilities and airport control towers.

The manner in which resources are allocated within the hierarchy depends on the planning time horizon, which can be broken down into tactical resource allocation, strategic resource allocation, and long-term airport slot allocation. These are each discussed below.

In the tactical time frame (less than two hours, and often minutes) resources are allocated by ATC in such a way as to maintain safety, with a particular emphasis on maintaining separation between aircraft. Resources such as airspace sectors, jet routes, arrival/departure routes, and airport runways are generally awarded on a first-come, first-served (FCFS) basis. This applies to ATC, where use of resources is decided by controllers, and tactical traffic flow management using tools such as the FAA's Time-Based Flow Management (TBFM) or EUROCONTROL's Enhanced Tactical Flow Management System (ETFMS). TBFM allocates arrival times at a network of metering points in order to maintain appropriate aircraft spacing given operating conditions. Allocated times are frozen after flights cross specified freeze horizons, allowing controllers to use speed control and vectoring to ensure aircraft meet these times. This type of demand management has particular applicability to UTM but would need to be adapted if a more federated architecture is adopted for UTM, with service providers self-separating their vehicles from other vehicles. Self-separation in ATM has been studied at length by NASA in the development of the Autonomous Operations Planner,[16] a flight-deck based suite of algorithms that detects and resolves conflicts with other aircraft while meeting route constraints.

Similar to TBFM, ETFMS arranges the flow of traffic through a regulated sector with the aim of maintaining the order of the flights. The resulting delays are allocated to the flights by delaying their departure times. In much the same way, TBFM allows local departures to an airport to be scheduled into an overhead stream of airborne arrival traffic to a given destination by delaying their departure time. Recent NASA work on improving the fairness of how this is done[17] may provide a useful model for pre-departure deconfliction in UTM for air taxi operations and pre-scheduled package deliveries.

FCFS in tactical allocation has long-standing acceptance in the aviation community, and generally makes sense for situations in which there may potentially be a physical queue for services. Different definitions of what constitutes "first-come" are likely to have very different implications for fairness, e.g., depending on whether "first-come" is defined by filing, pushing back, taking off, crossing a specified boundary, arriving, etc. Fairness issues associated with the timing of resource allocation are of particular relevance to UTM more broadly because of the proposed federated architecture. An approach similar to the freeze horizon used in TBFM may provide a solution, limiting how early resources can be requested, as suggested by Ref. 5 in the definition of a required time to act (RTTA), beyond which UTM flight plans are frozen.

In the strategic time frame – considered hours in advance of operation for traditional ATM – airspace and airport resources are allocated by air traffic flow management (ATFM). Resources are allocated in order to manage demand levels at those resources in such a way as to enable efficient separation management at the tactical level by controllers or tactical traffic management tools. Separation is not managed strategically because of the high uncertainties in conditions in the strategic time frame. In UTM, most operations will occur in the tactical time frame for traditional ATM, meaning that separation management may be possible pre-departure, with potentially reduced need for demand management. This will have to be confirmed for different operation types.

Both in the U.S. and in Europe there is heavy reliance on using scheduled times of operation as the basis for a claim to resources strategically. In structured ATFM programs, the basic paradigm is to create a virtual queue of airport or airspace "slots." In the U.S., allocation is based on the ration-by-schedule (RBS) algorithm[18] with earlier slots allocated to flights that were scheduled earlier. While most UTM operations will not be scheduled, a virtual queue based on some rationing scheme will likely have relevance to many types of operations.

RBS processes three nested queues of flights: scheduled flights; exempt flights; and flights that have received a prior allocation. In UTM, exemptions may be a mechanism to accommodate high-priority operations, such as public safety or emergency operations. In RBS, ties between exempt flights are broken by estimated time of arrival. In UTM, other criteria may be more appropriate.

RBS is executed in batch mode, and therefore also accounts for flights that have already received slots under a prior execution of RBS. This allows RBS to honor its prior allocations, rather than unfairly revoke previously awarded resources. This has similarities to the freezing of metering times by TBFM, which prevents changing resource allocation beyond a specified time. This could also be relevant to UTM, preventing the need for replanning close to departure.

In the U.S. ATFM system, RBS is supplemented by a dynamic compression algorithm that moves flights up in the arrival hierarchy to fill slots vacated by canceled flights, making more efficient use of arrival resources. Compression gives as much compensation as possible to the operator vacating the slot. Similar algorithms may be needed to accommodate canceled flights in UTM, which may be more common because of the on-demand nature of many UTM operations.

ATFM allocation practices and procedures have been heavily shaped in the U.S. by the joint FAA-industry venture known as Collaborative Decision Making (CDM),[18] which led to the creation of incentives for operators to submit timely and accurate information. Prior to CDM, airport arrival slot allocation was performed based purely on estimated time of arrival, which led operators to stop disclosing flight delays and cancellations to the FAA. CDM allowed operators to retain slots for canceled or delayed flights and use them for other operations. Similar mechanisms may be necessary in UTM to ensure truthful behavior.

Europe has taken a more authoritative stance on ATFM than the U.S. because of differences in the predictability of capacity constraints (impacted primarily by the lower incidence of convective weather). Airspace and airport allocation are applied days to weeks in advance and allow little or no overscheduling. Airport and airspace resources are allocated on a more proactive, continual basis than in the U.S., using a heuristic version of a large-scale optimization algorithm, in contrast to the as-needed batching approach used in the U.S. Such optimizations may be appropriate for a more centralized approach to UTM, or within single service providers in more federated architectures. In Europe, notification of assigned slots is not sent until 3 hours prior to expected gate pushback time.

In the U.S. and Europe, airport authorities, facilitated by trade bodies, impose long-term (i.e., months or years in advance) airport slot allocation on chronically congested airports. The processes rely heavily on slots that have been allocated historically, meaning preference is given to carriers that have operated for a qualifying duration. The allocation process takes place over an extended period of time, in which interactions and reviews take place with flight operators. Slots may be transferred or swapped between airlines or used as part of a shared operation. Slots may only be transferred to another airline that is serving or planning to serve the same airport.

The process is highly centralized and well established but contains significant subjective review and determinations by the coordinator. Efficient use of the resource is achieved simply by allocating all of the slots. The allocation of slots is based on a prioritization rule, heavily skewed toward incumbent operators and those operating a published schedule. This process necessarily assumes a lengthy petitioning, grievance, and allocation process (several months). To prevent awarded airport slots from going unused, utilization minima are set. Another key feature of the airport slot allocation process is that it guards against speculation, whereby slots are acquired merely to sell or manipulate pricing.

This type of highly procedural process will not work on the smaller time scales of real-time UTM (minutes or hours), nor is it clear what percent of operations at vertiports will be scheduled in advance. However, it has been proposed that some resources — even airspace blocks — could be allocated in advance by auction, without specifics of exactly how they will be used (Skorup, 2018).[19] However, because access to the airspace is based entirely on the criteria for winning the auction, and only one operator gains that right to the exclusion of all other operators, it would be very difficult to ensure fair access to resources for all users in such a scheme.

Non-Aviation Domains

Fairness in the allocation of resources has also been considered across a wide range of domains outside aviation, such as:

- Non-discriminatory train dispatching in the rail transport market;[20]
- Wireless networks, where radio spectrum is allocated in real time;[21,22]
- Automated trial decisions, where bias must be avoided in judicial decision making;[23]
- Free museum ticket allocation, which allocates free timed-entry passes at high-demand times;
- Legislative seat assignments in the U.S. Congress;
- Estate inheritance;
- Supplying organs to transplant patients, which uses a points-based schema, based on weighting a number of different factors;[24]
- Army discharges at the end of WWII, which also used a points-based schema.[24]

There are several valuable lessons to learn from approaches to resource allocation in these various fields. One lesson learned is that a structured, points-based schema (first rights go to the claimant with the most points accumulated) can be an effective way to prioritize for services. This is usually more palatable than relative ranking (claimant *A* always has priority over claimant *B*). Prioritization in the cellular services industry has demonstrated that structured procedures can perform rapidly and are capable of balancing efficiency with fairness, although these procedures may not translate well to UTM. Information packets are mostly uniform, and the consequences of dropping or rerouting them are less severe. Statistical inference techniques (machine learning) can be applied to track performance over long periods of time and correct for bias.

Another lesson learned is that market-based mechanisms can be an effective way to alleviate all or most of the need for centralized allocation, although some regulation is often still required. Market-based mechanisms are desirable because they are often the best way to ensure that each stakeholder's welfare is represented and maximized to the extent possible. Market-based solutions do, however, optimize efficiency rather than fairness, so have to be designed with fairness in mind. Auctions are a common way to implement market-based mechanisms. Recent technologies, such as online bidding by advertisers wanting to post ads to online viewers, has shown that electronic auctions can be conducted at high speed without the need for human intervention. Before applying this to a UTM environment, however, one would have to consider the technological sophistication required of all users. Congestion pricing is a common alternative to auctions, but it suffers from the defect that there is usually no basis for setting the appropriate prices.

Perhaps the most important lesson learned is that air transportation is orders of magnitude more complex than the commonly cited resource allocation situations outside aviation. Air transportation has technical, stochastic, political, legal, dynamic, and safety aspects that must be considered. Moreover, resource allocation techniques necessarily vary by the nature of the resources and the claimants. For both these reasons, allocation paradigms ported over from other domains may require significant modification.

Unmanned Traffic Management

While the topic of fairness has not been explored in depth within the UTM ecosystem, a number of works have examined emergent behavior of unmanned traffic in complex and highly utilized airspace.[25,26,27,28] These works, however, do not consider varying operator behavior that could lead to questions of fair airspace usage, particularly in the context of multiple interacting operations and deconfliction. A number of systems have been proposed and evaluated that rely on a centralized entity to coordinate unmanned traffic[29] or on a well defined airspace structure that organizes unmanned traffic with similar operational objectives.[30] While these approaches consider operations at scale, they do not examine how multiple decision makers with unique operational preferences could coexist.

The digitized and autonomous nature of UTM make it a good candidate for application of methodologies that have been proposed for ATM, but have not found their way to implementation, like automated, negotiation based conflict resolution.[31] Other methods developed for agent-based automated resource allocation and negotiation[32] are potential candidates for implementation in UTM. However, a number of challenges exist in applying these methods to the aviation domain, where the safety-critical requirements and the complex distributed nature of the system can make these methods impractical in implementation. Questions about fair airspace access arise as well, but little work has been done in analyzing these questions quantitatively.

Overall, an analysis gap exists for autonomous coordination and negotiation in UTM, which this work aims to fill through quantitative examination of fairness in the ecosystem. Approaches for fair resource allocation in ATM, which are based on FCFS allocation, could be appropriate to UTM, but may require modification to ensure fairness in the more federated architecture proposed for UTM. Lessons can also be learned from resource allocation in other domains, but because these systems are generally orders of magnitude less complex than ATM, significant modification may be required. This paper therefore seeks to explore the fairness implications of a FCFS allocation in UTM through simulation.

III. Objectives

As described in Section II, FCFS is a widely accepted approach to resource allocation in ATM, both tactically and strategically. It is also an approach that can be applied in the more federated architectures proposed for UTM. The objectives of this paper are therefore as follows:

1. Explore the fairness implications of applying a sample FCFS algorithm to the allocation of airspace access in UTM by estimating metrics describing fairness for sample UTM traffic scenarios; and
2. Identify key considerations that must be dealt with to ensure fair allocation by FCFS algorithms in UTM.

IV. Experimental Setup

Simulation is used to estimate metrics describing fairness across operators for two sample UTM traffic scenarios. The allocation algorithm is described below, followed by a description of the scenarios simulated and the experimental design matrix. This is followed by a description of the metrics used to describe fairness, and the simulation environment in which the experiments are completed.

Allocation Algorithm

The sample FCFS allocation considered in this paper is for strategic deconfliction of flight plans (i.e., pre-departure), and is based on when the operator filed the flight plan. These scenarios therefore assume a simple first-requested, first-served (FRFS) allocation method. Under this scheme, flight plans filed first – regardless of their requested departure times – have priority over flight plans filed later. In the case of conflict, there is assumed to be no direct USS-to-USS negotiation: the operator who files first always takes priority, and retains their flight plan, while the later-filing operator must adjust their flight plan to deconflict.

We hypothesize that in this type of scenario, operators will try to minimize costs associated with deconflicting with higher priority flights by filing as far in advance of departure time as possible, increasing their odds of a direct route. In real operations, how far each operator is able to file in advance will be limited by factors specific to that operator.

Scenarios Simulated

Two sample UTM traffic scenarios are considered:

1. A package delivery scenario with two operators serving overlapping regions, and
2. An air taxi scenario with two operators serving a network of vertiports.

In the package delivery scenario, flights for two operators are simulated serving an overlapping region from separate warehouses, as shown in Figure 1a. Because both operators serve the same region, there is significant interaction between flights from different operators. This requires that all flights be checked for conflicts with the other operators' flights, and any conflicts be resolved. All flights originate at the operator's warehouse, while the delivery sites (flight destinations) are randomly distributed across the specified region according to a Gaussian distribution with the standard deviation forming a circle with radius of 6km. For simplicity, only the outbound flight segments, from the warehouse to the delivery site, are simulated in both variations. The return segment is not simulated. Other simulation parameter values are presented in Table 1.

A single-operator package delivery case is also simulated in which one operator serves a region surrounding its warehouse, as shown in Figure 1b. Because there is only one operator, there is no need for the operator to deconflict with flights from other operators, although there is still a need to deconflict with its own flights. In this case the operator is therefore operating without any impact from other operators.

In the air taxi scenario, flights for two operators are simulated serving a network of seven vertiports, located at the vertices and center of a hexagon (with vertices located 16km from the center), as shown in Figure 2a. Both operators serve the same network of vertiports, and also share the vertiports themselves, so there is significant interaction between flights from different operators. This requires that all flights be checked for conflicts with the other operators' flights, and any conflicts be resolved. Other simulation parameter values are presented in Table 1.

A single-operator air taxi case is also simulated in which one operator serves the entire network, as shown in Figure 2b. Because there is only one operator, there is no need for the operator to deconflict with flights from other operators,

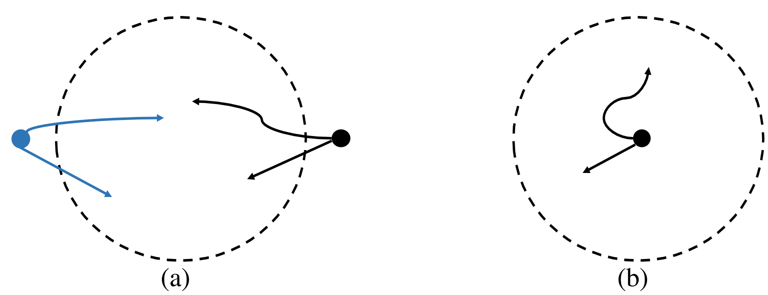

Figure 1. Notional package delivery example with (a) two operators (blue, black) delivering from separate warehouses (blue and black dots) 12km apart, to serve demand in a common region (defined by dashed lines); and (b) a single operator serving a region surrounding its warehouse. In both cases the delivery regions are defined by Gaussian distributions with standard deviation forming a circle of radius 6 km.

Table 1. Simulation parameter values

	Package delivery	*Air taxi*
Number of operators	[2; 1]	[2, 1]
Origin type	Point source	Point source
Number of origins	2	7
Destination type	Area sink (Gaussian distribution)	Point sink
Number of destinations	(unlimited)	7
Average source demand per operator	[25; 50; 100; 250] requests/hr	[25; 50; 100; 250] requests/hr
Source demand distribution	Poisson	Poisson
Cruise speed	20 ms^{-1}	60 ms^{-1}
Cruise altitude	120 m	[500, 800] m
Separation requirement	100 m	150 m
File-ahead time: Operator 1	[0; 100; 250; 450; 900; 1,800] seconds	[0; 100; 250; 450; 900; 1,800] seconds
File-ahead time: Operator 2	0 seconds	0 seconds

although there is still a need to deconflict with its own flights. In this case the operator is therefore operating without any impact from other operators.

In both the package delivery and air taxi scenarios demand is generated randomly for each source, based on a Poisson distribution with specified average rate per hour shown in Table 1. In the package delivery scenario, there are two sources – one at each operator's warehouse. In the air taxi scenario, there are seven sources, located at each vertiport. A range of demand rates are simulated in each case representing low, medium and high traffic demand, shown in Table 1. Simulation durations vary from 6 hours to 120 hours (in simulation time), depending on the demand rates simulated. These values are sufficiently long to ensure steady state conditions. Flights for the first 30 minutes of each simulation are discarded so as to prevent skewing the results by the demand ramp up period.

In all scenarios simulated in this paper, conflicts are identified by comparing full 4-dimensional trajectories (4-DTs). The NASA UTM concept specifies flight intent using flight volumes,[6] which could extend well outside the flight's planned 4DT. However, deconfliction of 4DTs makes the most efficient use of the airspace, and is therefore most conservative in terms of delay allocation. For this reason, flight plan intent for all scenarios simulated in this paper is specified in the form of a 4DT. Flights are deconflicted by adjusting the departure time only. No changes are made to the flight plan routing. While such an approach is not always the most efficient, it simplifies the analysis as all delay is incurred on the ground, pre-departure, with flight times remaining unchanged relative the requested flight plans. This solution to deconfliction can cause very high delays for flights in conflict operating in opposite direction between two points. For this reason, we also introduce procedural separation between flights in the air taxi scenarios, with flight altitude based on flight direction: flights operating on routes with a northbound heading cruise at an altitude of 500m, while flights operating on routes with a southbound heading cruise at an altitude of 800m. Conflicts during approach and descent to/from these altitudes are still addressed.

In the baseline simulations for the cases in which operators are competing for resources (shown in Figure 1a and Figure 2a), both operators are assumed to request service by filing their flight plans at the time when they plan to depart. This is therefore truly on-demand service. In order to explore the impact of file-ahead time on fairness and

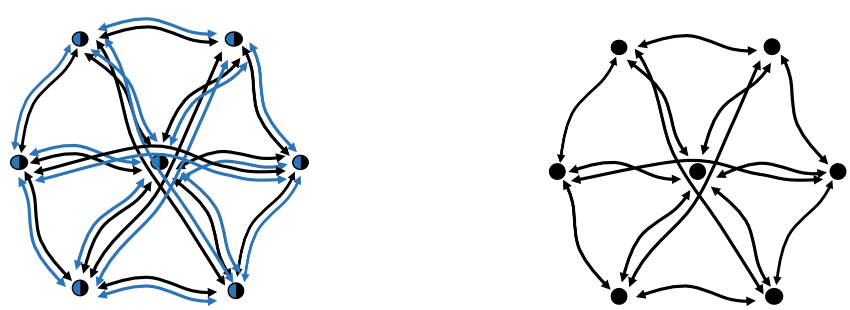

(a) (b)

Figure 2. Notional air taxi example for a network of 7 vertiports served by (a) two operators (blue, black) and (b) a single operator. In both cases the center vertiport is located equidistant from the other six vertiports at a distance of 16km.

other metrics, other file-ahead scenarios are simulated (listed in Table 1), in which one operator is simulated filing all flight plans a fixed time before planned departure time, while the other operator still files at their planned departure times. The operator that files ahead gains an advantage by filing early, because of the FRFS prioritization based on when flight plans are filed. In reality, different types of operations would have different flexibility in when they could file.

An experimental design matrix for the simulations carried out for this paper is shown in Table 2. Simulation durations are also shown in this table for completeness.

Metrics for Fairness

Many metrics proposed to quantify fairness are based on concrete notions of claims and entitlement associated with resources. However, because neither the resources nor the claims to them have been clearly established in UTM, a different approach is required which measures the degree to which different participants have been denied free access to the airspace. For example, for traditional ATM, Ref. 33 and 34 propose a metric for whether aircraft are receiving even treatment of ATC separation services which is strongly tied to delay costs. The distribution of costs across operators can then be evaluated to quantify fairness. Costs accounting for operator utility can be considered if such data is available, or they can be represented by metrics for each operator that are known to the system.

For this paper, we quantify fairness by comparing the distribution of costs across operators using the normalized fairness metric in Equation 1.[34] Operator cost is defined as a function of average ground delay in Equation 2.

Equation 1 defines the normalized fairness metric E as the ratio between the geometric and arithmetic mean of operator costs P_i across operators i.

$$E = \frac{\left(\prod_{i=1}^{n} P_i + \varepsilon\right)^{1/n}}{\sum_{i=1}^{n} P_i + \varepsilon} \cdot n \tag{1}$$

P_i is defined by a penalty function that quantifies the cost of delay to the operator in Equation 2. n represents the number of operators, and ε is a small, strictly positive number used to ensure that when the average ground delay is zero for one operator of the set of n operators, the normalized fairness metric E does not result in zero. The normalized fairness metric E has a value between 0 (unfair) and 1 (fair).

The cost of delay to operators is unlikely to be linear with average ground delay D_i. Operators are likely to be indifferent to very low delays, but intolerant to very high delays. We therefore define a piecewise penalty function based on two delay thresholds, $D_{indifferent}$ and $D_{intollerable}$, as follows:

$$P_i(D_i) = \begin{cases} \sqrt{D_i} & D_i < D_{indifferent} \\ D_i + \alpha & D_{indifferent} \le D_i \le D_{intollerable} \\ {D_i}^2 + \beta & D_i > D_{intollerable} \end{cases} \tag{2}$$

The exact functions and threshold values used in the piecewise penalty function are likely to be highly dependent on the operator. For this paper, we use the functions shown in Equation 2, and thresholds of 1 minute and 5 minutes for $D_{indifferent}$ and $D_{intollerable}$, respectively. Both values are applied to both operators in both the package delivery and air taxi scenarios. α and β are defined to ensure that the piecewise function is continuous at $D_{indifferent}$ and $D_{intollerable}$, as follows:

$$\alpha = \sqrt{D_{indifferent}} - D_{indifferent} \tag{3}$$

$$\beta = D_{intollerable} + \sqrt{D_{indifferent}} - D_{indifferent} - {D_{intollerable}}^2 \tag{4}$$

Other values and functions could be applied to tailor the fairness metric to a particular application or operator preference.

Simulation Capability

A high-fidelity airspace simulator is used to model traffic in each scenario. In the simulation, flight requests are generated using a stochastic demand generation process, which specifies the requested departure time based on the file-ahead time for the scenario being run. The requests are passed on to a planner which calculates the ground delay required to deconflict the flight plan with all other flight plans already requested. To model the vehicles in the simulation, we use a simple point-particle dynamic model with a hybrid proportional-integral-derivative (PID) and

Table 2. Experimental Design Matrix

Simulation No.	*Operation Type*	*Number of Operators*	*Hourly Demand Rate [ac/hr]*	*Simulation Time [hours]*	*File-ahead time Operator 1[1] [s]*
1	Package delivery	1	25	120	0
2			50	60	0
3			100	30	0
4			250	12	0
5		2	25	120	0
6					100
7					250
8					450
9					900
10					1,800
11			50	60	0
12					100
13					250
14					450
15					900
16					1,800
17			100	30	0
18					100
19					250
20					450
21					900
22					1,800
23			250	12	0
24					100
25					250
26					450
27					900
28					1,800
29	Air taxi	1	25	60	0
30			50	30	0
31			100	15	0
32			250	6	0
33		2	25	60	0
34					100
35					250
36					450
37					900
38					1,800
39			50	30	0
40					100
41					250
42					450
43					900
44					1,800
45			100	45	0
46					100
47					250
48					450
49					900
50					1,800
51			250	18	0
52					100
53					250
54					450
55					900
56					1,800

logic control for guidance. While the simulated vehicles have on board sensing and conflict resolution capabilities, we do not consider them in this work.

In all cases, individual trajectories are optimized to maximize a cost function as described in Ref. 35. In this paper, the trajectory optimization does not consider other traffic – flights are deconflicted strategically through departure time only using a scheduling service.

Given a requested plan and all the previously approved flight-plans, the scheduling service finds potential conflicts with the requested plan and uses a linear programming optimization to compute the minimum ground delay for the request such that it does not conflict with previously approved plans. Note that even though a previously approved plan may have a departure time in the future, the optimization does not discard it since this flight could potentially conflict with the requested flight if the requested plan is sufficiently delayed. This is important to consider especially for scenarios with high traffic demand, where flights may be significantly delayed before potential conflicts get resolved.

V. Simulation Results

Simulation results are presented below for the two operation types simulated:

1. A package delivery scenario with two operators serving overlapping regions, and
2. An air taxi scenario with two operators serving a network of vertiports.

Simulating Package Delivery Operations

Figure 3 shows the results from the package delivery simulations. Comparing across subplots (a) to (d), the average ground delay for the single operator baseline (black dashed lines in Figure 3) increases with increasing demand, as expected – from under 1 second for demand of 25 requests/hr per operator (Figure 3a) to approximately 6 minutes for demand of 250 requests/hr per operator (Figure 3d). Note that the scale of the vertical axis varies significantly across subplots. Figure 3 also shows the average simulated ground delay for two operators (the blue and red lines), with one operator (operator 1) filing ahead by between 0 and 30 minutes (1,800 seconds), across the same range of demand levels. Error bands showing one standard deviation of simulated ground delay are also shown.

With both operators filing at the desired departure time (0 file-ahead time), the average simulated ground delays are significantly higher than for the single operator at all demand levels – because of the doubling of traffic caused by the additional operator. In all cases, both operators experience similar delays, only varying because of the stochastic nature of the demand. As the file-ahead time for operator 1 increases, its average ground delay decreases (red dashed line in Figure 3), while the average ground delay of operator 2, that still files at the desired departure time, increases (blue solid line in Figure 3). The ability of operator 1 to file-ahead therefore has a clear impact not only on its own average delay, but also on that of its competitor. The effect is similar across all demand levels, although there are differences in the rates at which the average delays change, and, more significantly, on the magnitude of the delay.

As the file-ahead time of operator 1 increases, its average ground delay approaches that of the single-operator baseline at all demand levels. However, the file-ahead time at which the operator 1 delay approximates the single-operator delay varies, from 450 seconds (7.5 minutes) at demand of 25 requests/hr, to approximately 900 seconds (15 minutes) at demand of 100 requests/hr. This indicates that at these file-ahead times, operator 1 operates almost as if it were the only operator in the airspace. At 250 requests/hr, operator 1 delay never approximates the single operator delay, even at the highest file-ahead time simulated of 1,800 seconds (30 minutes). This suggests that higher file-ahead times would be needed for operator 1 to operate as if it were the only operator in the airspace.

The rate at which the average ground delay of operator 2 increases slows with increasing operator 1 file-ahead time, ultimately plateauing to a maximum value for demand rates of 25 requests/hr and 50 requests/hr (Figure 3a and b). Higher file-ahead times appear to be necessary to reach such a maximum value for the demand rates of 100 requests/hr and 250 requests/hr (Figure 3c and d).

The key difference between the subplots in Figure 3 is the magnitude of the simulated delays. At low demand levels of 25 requests/hr (Figure 3a), average ground delays and the difference between operator 1 and 2's delay never exceeds 9s. At high demand levels of 250 requests/hr (Figure 3d), however, the average ground delay for operator 2 exceeds 1 hour 45 minutes when operator 1 files ahead by 30 minutes while operator 1's average ground delay is just over 30 minutes. Fairness is far more likely to be considered an issue when the magnitude of the delays, and delay difference between operators, is high. This effect is captured by the normalized fairness metric, shown in Figure 4.

As expected, Figure 4 shows a normalized fairness metric of 1.0 at all demand levels when the operator 1 file-ahead time is 0 seconds, since no operator gains any advantage by filing early. As the operator 1 file-ahead time increases, the normalized fairness metric decreases. For demand values of 50 requests/hr and 100 requests/hr, this decrease in fairness is small – under 15%, even for an operator 1 file-ahead time of 30 minutes. This is because the

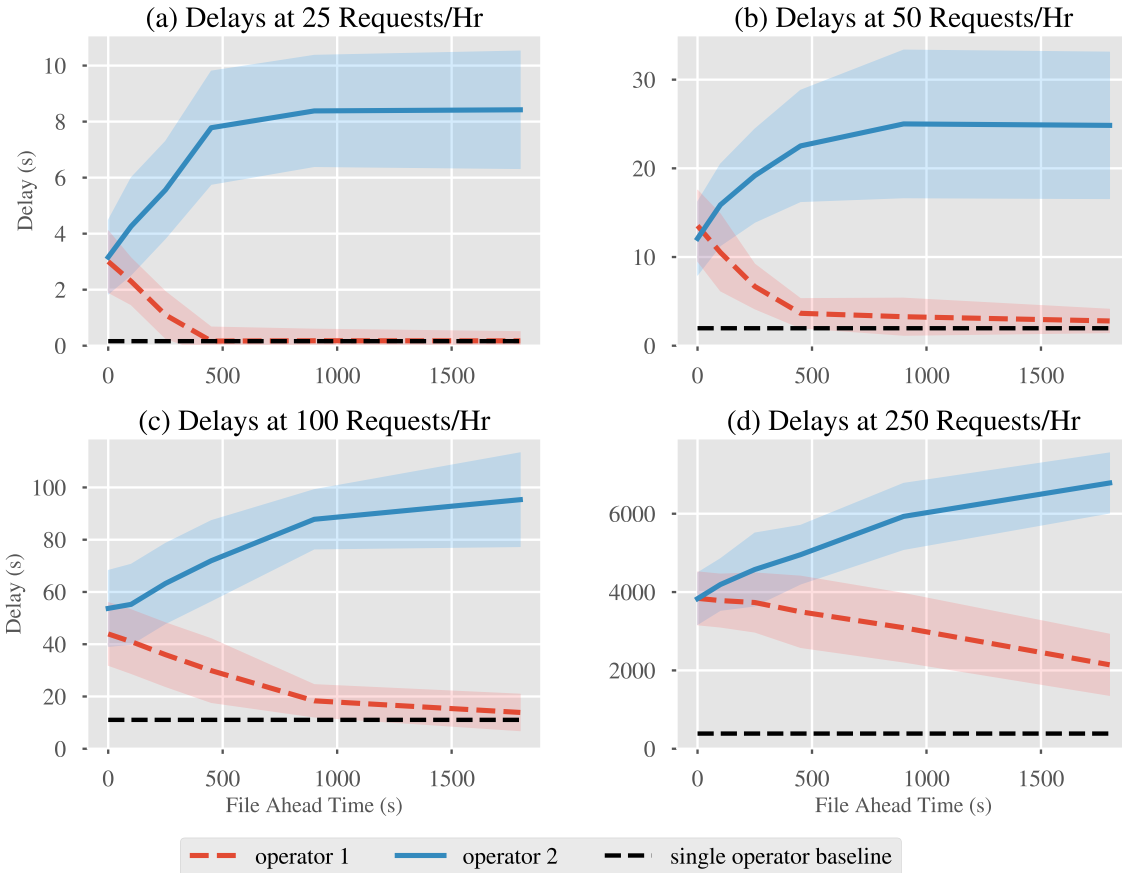

Figure 3. Simulated average ground delay for two package delivery operators, under varying operator 1 file-ahead times, for demand per operator of (a) 25 requests/hr; (b) 50 requests/hr; (c) 100 requests/hr; (d) 250 requests/hr. In all cases, the file-head time for operator 2 is 0 seconds. Simulated ground delays for a single operator baseline are also shown (black dashed lines) for reference. Note that that scales on the vertical axis vary.

magnitude of the average ground delay for operator 2 is small – under 2 minutes – even if it is still significantly larger than the average ground delay of operator 1. This is no longer the case when demand increases to 250 requests/hr, with operator 2's average ground delay increasing to over 1 hour 45 minutes when operator 1 files-ahead by 30 minutes. This results in a lower normalized fairness metric of under 0.6 as the high delay of operator 2 is penalized by the penalty function from Equation 2. The normalized fairness metric at this high demand level only reduces significantly at the higher operator 1 file-ahead times – above 250 seconds.

The normalized fairness metric at very low demand (25 requests/hr) also decreases significantly with increasing operator 1 file-ahead time – by up to 30% when operator 1 files-ahead by more than 250 seconds – despite the very low magnitude of delay in the order of seconds at 25 requests/hr. The reason for this is the very low delay for operator 1 at these demand levels, which is close to zero at the higher file-ahead times. This reduces the numerator of Equation 1, resulting in the low normalized fairness metric despite the square root in the penalty function (Equation 2). Other penalty functions could be identified to mitigate this effect.

It is noted that the normalized fairness metric is fairly sensitive to the penalty function chosen, impacting particularly this latter effect of low fairness values at low demand levels. Using a linear penalty function equal to delay, the normalized fairness metric with a 30-minute file-ahead time is between 0.35 (at 25 requests per hour) and 0.85 (250 requests per hour).

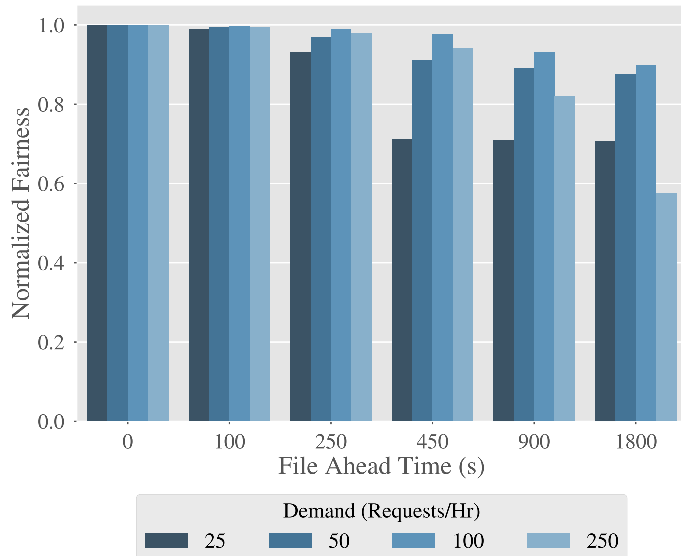

Figure 4. Simulated normalized fairness metric, calculated using Equation 1, for two air package delivery operators under varying demand per operator and operator 1 file-ahead time. Operator 2 files at the desired departure time (file-ahead time = 0).

Simulating Air Taxi Operations

Figure 5 shows the results from the air taxi simulations. The average ground delay for the single operator baseline (black dashed lines in Figure 5) again increases with demand, as expected, from under 1 second for demand of 25 requests/hr (Figure 5a) to 43 seconds for demand of 250 requests/hr (Figure 5d).

With both operators filing at the desired departure time (0 file ahead time), the average simulated ground delays are again significantly higher than for the single operator at all demand levels – because of the doubling of traffic caused by the additional operator. As the file-ahead time for operator 1 increases, its average ground delay decreases, while the average ground delay of operator 2, that still files at the desired departure time, increases. The ability of operator 1 to file-ahead therefore also has a clear impact not only on its own average delay, but also on that of its competitor, like in the package delivery scenarios. The effect is again similar across all demand levels, although there are differences in the rates at which the average delays change and the magnitude of the delay.

As the file-ahead time of operator 1 increases, its average ground delay approaches that of the single-operator baseline at all demand levels. The rate at which the average ground delay of operator 2 increases also slow, plateauing to a maximum value for demand rates of 100 requests/hr and below. Higher file-ahead times appear to be necessary to reach such a maximum value for the demand rate of 250 requests/hr (Figure 5d).

As for the package delivery scenarios, the key difference between the subplots in Figure 5 is the magnitude of the simulated delays. At low demand levels of 25 requests/hr (Figure 5a), average ground delays and the difference between operator 1 and 2's delay never exceeds 7s. At high demand levels of 250 requests/hr (Figure 5d), however, the average ground delay for operator 2 exceeds 20 minutes when operator 1 files ahead by 30 minutes while operator 1's average ground delay is approximately 1 minute.

The magnitude of average ground delay simulated in the package delivery scenarios (Figure 3) is consistently higher than in the air taxi scenarios (Figure 5), and particularly at 250 requests/hr per operator, when the maximum package delivery delay of over 2 hours compares to the maximum air taxi delay of 27 minutes. The reason for this is thought to be that the terminal airspace capacity around a single warehouse in the package delivery scenarios is significantly smaller than in the 7 vertiport network for the air taxi scenarios, leading to significantly higher delays when similar separation requirements are enforced. A modified terminal area separation policy to allow higher volumes of operations from a single package delivery warehouse could mitigate these high delays.

Figure 6 shows the normalized fairness metrics for the air taxi scenarios. As the operator 1 file-ahead time increases, the normalized fairness metric decreases. For demand values of 50 requests/hr and 100 requests/hr, this decrease in fairness is small – up to 8%, even for an operator 1 file-ahead time of 30 minutes. This is because the magnitude of the average ground delay for operator 2 is small – under 1 minute – even if it is still significantly larger

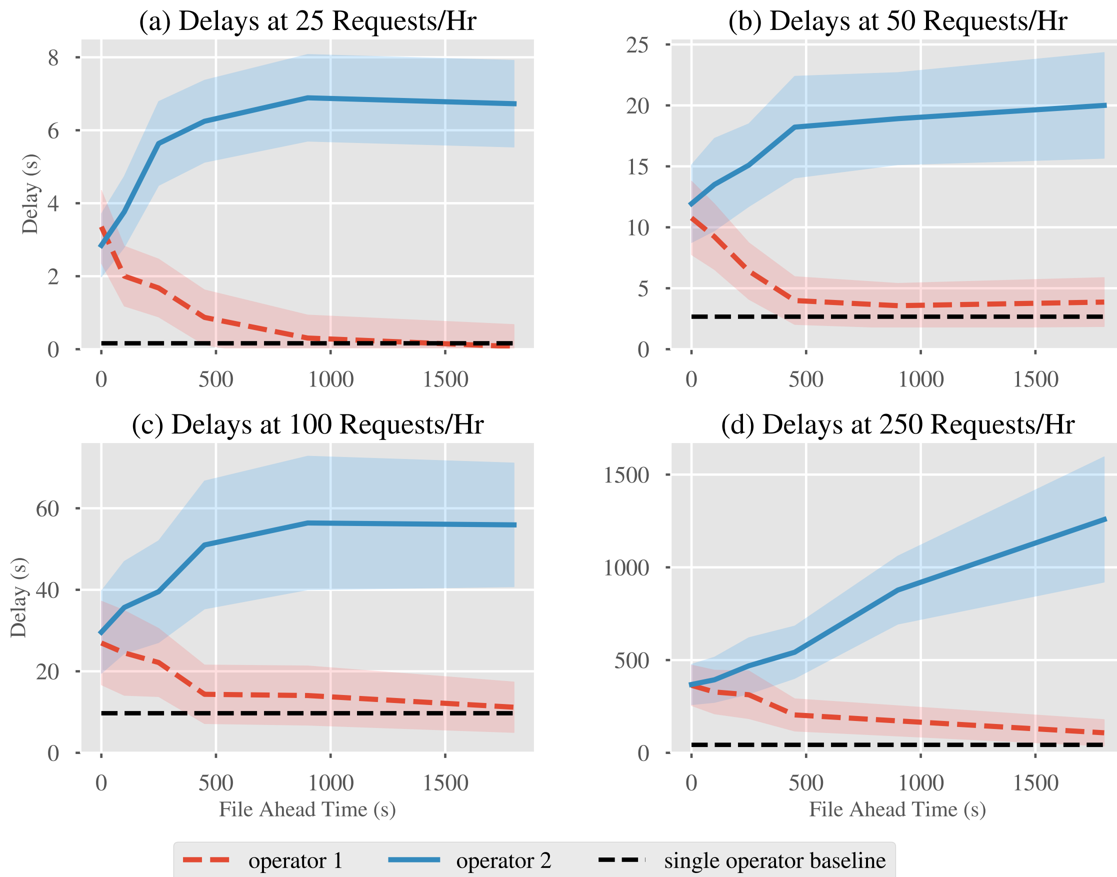

Figure 5. Simulated average ground delay for two air taxi operators, under varying operator 1 file-ahead times, for demand per operator of (a) 25 requests/hr; (b) 50 requests/hr; (c) 100 requests/hr; (d) 250 requests/hr. In all cases, the file-head time for operator 2 is 0 seconds. Simulated ground delays for a single operator baseline are also shown (black dashed lines) for reference. Note that that scales on the vertical axis vary.

than the average ground delay of operator 1. This is no longer the case when demand increases to 250 requests/hr, with operator 2's average ground delay increasing to over 20 minutes when operator 1 files-ahead by 30 minutes. This results in a very low normalized fairness metric of under 0.2 – significantly lower than any of the normalized fairness metrics for the package delivery scenarios (Figure 4). Like the package delivery scenarios though, the normalized fairness metric only reduces significantly above 250 seconds.

The normalized fairness metric at very low demand (25 requests/hr) also decreases with increasing operator 1 file-ahead time – by up to 35% when operator 1 files-ahead by 30 minutes. While not as extreme as for high demand, it is more extreme than for demand levels of 50 and 100 requests/hr, despite the very low magnitude of delay at 25 requests/hr. Like for the package delivery scenarios, the reason for this is the very low delay for operator 1 at these demand levels, which is essentially zero at the higher file-ahead times.

VI. Conclusions

This paper explores the implications of different approaches to strategic decentralized trajectory deconfliction in UTM on fairness. Lessons learned from approaches to allocate resources in traditional ATM and other domains are highlighted. Points- and market-based schema have proven to be effective in other domains, but these domains are generally significantly less complex than ATM. FCFS allocation of resources has long-standing acceptance in the aviation community. It is applied tactically by air traffic controllers and tactical traffic flow management tools such as the FAA's TBFM or EUROCONTROL's ETFMS, and strategically by ATFM in the form of virtual queues of

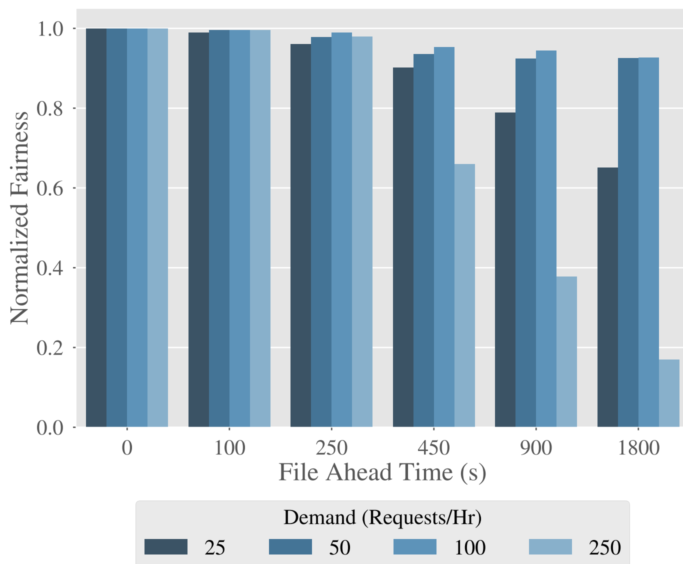

Figure 6. Simulated normalized fairness metric, calculated using Equation 1, for two air taxi operators under varying demand per operator and operator 1 file-ahead time. Operator 2 files at the desired departure time (file-ahead time = 0).

airport or airspace 'slots' allocated according to scheduled times of operation. However, it is unclear what the fairness implications are for a FCFS allocation of access to airspace in UTM.

In this paper, simulation is used to provide quantitative examples to explore how a FCFS approach to strategic deconfliction – based on when operators file their flight plans – performs in terms of fairness, quantified by comparing average ground delay across operators and by calculating a normalized fairness metric that accounts for operator cost of delay. Two scenario types are simulated: two package delivery operators serving a common region from separate warehouses; and two air taxi operators serving the same network of 7 vertiports.

For the simulated scenarios, results suggest that if a package delivery operator is able to file ahead by as much as 30-minutes under high demand conditions (250 requests/hr per operator), the average ground delay incurred due to strategic deconfliction, based on a FRFS allocation between two operators, could be 70% lower than for an operator that was not able to file ahead. The average ground delay of that latter operator could be as high as 1 hour 45 minutes. For air taxi services between shared vertiports under similar file-ahead and demand conditions, average ground delay for the operator filing early could be 95% lower than for the operator that is not able to file ahead, while the average ground delay of that latter operator could be as high as 20 minutes. These results are, however, highly sensitive to the demand levels simulated. A comparison of a normalized fairness metric across the different scenarios shows that, for the scenarios simulated, it is primarily at the high demand levels and file-ahead times greater than 250 seconds that fairness becomes a critical issue. These results highlight that there could be significant inequity if the file-ahead time is not constrained in a FRFS allocation, but primarily at high demand levels.

Such a constraint on file-ahead time could be applied in a similar way to the freeze horizon in TBFM, or as described in Ref. 5 in the form of a Required-Time-To-Act (RTTA). This is a time before flight operation after which it becomes difficult to change the requested flight plan. Under this concept, all flights would be considered to have equal priority before their RTTAs, with the exception of specific high priority operations (e.g., emergency services). Any conflict resolution prior to the RTTA would not be governed by FCFS or FRFS rules. After the RTTA, however, previously filed flight plans would take priority over later filed flight plans – effectively being 'frozen'. These frozen flight plans would be protected from any further change in all but the most extreme situations.[5] Ramifications of this approach, however, would have to be studied carefully.

The simulation results presented in this paper represent the most extreme imbalance in departures delays possibly under different RTTA values for the simulated scenarios. With the file-ahead time of operator 1 defining the RTTA, all of operator 1's flight plans are 'frozen' at their RTTAs, while operator 2 only files its flight plans at its desired departure times. Using the normalized fairness metric described in Equation 1, and a value for this metric that is considered acceptable, the results presented in this paper suggest that RTTA values can be identified for different types of operations. For both the package delivery and air taxi scenarios simulated, an RTTA of 250 seconds would provide normalized fairness metrics of at worst 0.95. These results are however, sensitive to the penalty function used to define the normalized fairness metric, described in Equation 2.

Future work will explore how the quantitative results outlined in this paper can be tied to a theoretical foundation. This will allow the work presented here to be both more generalizable, and to provide guarantees on the expectations of certain outcomes in the UTM design space relating to fairness. In particular, a proof or lack thereof of Pareto Optimality for protocols in decentralized strategic deconfliction would be an important part of understanding and designing systems governing the future of UTM. Many other approaches also exist for the allocation of resources in UTM that do not rely on FCFS, including not requiring the resolution of conflicts strategically, or optimizing the allocation of delay strategically across all operators by a centralized deconfliction service. The analysis that explores the fairness and efficiency implications of these approaches is left for future work.

References

[1] Balakrishnan, K., Polastre, J., Mooberry, J., Golding, R., and Sachs, P., "Blueprint for the sky: The roadmap for the safe integration of autonomous aircraft," *Airbus UTM*, San Francisco, CA, 2018.

[2] Jenkins, D., Vasigh, B., Oster, C., and Larsen, T., "Forecast of the Commercial UAS Package Delivery Market," *Embry-Riddle Aeronautical University*, 2017.

[3] Booz Allen Hamilton, "Urban air mobility market study," Presentation to NASA Aeronautics Research Mission Directorate, 2018, URL: https://go.nasa.gov/2MVSbth [retrieved 27 November 2019].

[4] Kopardekar P., Rios, J., Prevot, T., Johnson, M., Jung, J., and Robinson. J., "Unmanned aircraft system traffic management (utm) concept of operations," *AIAA Aviation Forum*, 2016.

[5] Hately, A., *et al.*, "*CORUS U-Space Concept of Operations*," *2016 SESAR 2020 RPAS Exploratory Research Call*, EUROCONTROL, 2019.

[6] Federal Aviation Administration, "Unmanned Aircraft System (UAS) Traffic Management (UTM) Concept of Operations v1.0," *FAA NextGen Office*, 2018.

[7] Chan, W.N., Barmore, B., Kibler, J., Lee, P.U., O'Connor, N., Palopo, K., Thipphavong, D.P. and Zelinski, S., "Overview of NASA's ATM-X Project," In *2018 Aviation Technology, Integration, and Operations Conference*, 2018, p. 3363.

[8] Dukowitz, Z. "Switzerland First Country with Nationwide UTM to Launch Medical Drone Delivery Program," *UAV Coach* [online news report], URL: https://uavcoach.com/switzerland-utm/ [retrieved 5 December 2019].

[9] Unmanned Airspace, "December FAA and ASTM remote ID publications "will transform commercial drone operations worldwide," *Unmanned Airspac*e [online news report], URL: https://www.unmannedairspace.info/emerging-regulations/december-faa-and-astm-remote-id-publications-could-transform-global-commercial-drone-operations/ [retrieved 5 December 2019].

[10] Rios, J., Mulfinger, D., Homola, J. and Venkatesan, P., "NASA UAS traffic management national campaign: Operations across Six UAS Test Sites." In *2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)*, 2016, pp. 1-6.

[11] Federal Aviation Administration, "Package Delivery by Drone (Part 135)," *Federal Aviation Administration*, 2019 URL: https://www.faa.gov/uas/advanced_operations/package_delivery_drone/ [retrieved 27 November 2019].

[12] ICAO, "Unmanned Aircraft Systems Traffic Management (UTM) – A Common Framework with Core Principles for Global Harmonization," *ICAO*, 2019, URL: https://www.icao.int/safety/UA/Documents/UTM-Framework.en.alltext.pdf [retrieved 27 November 2019]

[13] Rios, J., "Strategic Deconfliction: System Requirements", *NASA UAS Traffic Management (UTM) Project*, July, 2018, URL: https://utm.arc.nasa.gov/docs/2018-UTM-Strategic-Deconfliction-Final-Report.pdf [retrieved 27 November 2019].

[14] O'Neill, B., "A Problem of Rights Arbitration from the Talmud," *Mathematical Social Sciences* 2:345–371, 1982.

[15] Aumann, R.J., and Maschler, M., "Game Theoretic Analysis of a Bankruptcy Problem from the Talmud," *Journal of Economic Theory* 36:195–213, 1985.

[16] Karr D, Vivona R, Roscoe D, Depascale S, Wing D, "Autonomous operations planner: A flexible platform for research in flight-deck support for airborne self-separation," In *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2012, p. 5417.

[17] Smith, N.M., Brasil, C., Lee, P.U., Buckley, N., Gabriel, C., Mohlenbrink, C.P., Omar, F., Parke, B., Speridakos, C. and Yoo, H.S., "Integrated demand management: Coordinating strategic and tactical flow scheduling operations", *16th AIAA Aviation Technology, Integration, and Operations Conference*, 2016, p. 4221.

[18] Wambsganss, M.C., "Collaborative decision making in air traffic management," In *New Concepts and Methods in Air Traffic Management* (pp. 1-15). Springer, Berlin, Heidelberg, 2001.

[19] Skorup, B., "Auctioning Airspace," Working Paper, *Mercatus Center at George Mason University*, Arlington, Va., November 2018.
[20] Luan, X., Corman, F. and Meng, L., "Non-discriminatory train dispatching in a rail transport market with multiple competing and collaborative train operating companies," *Transportation Research Part C: Emerging Technologies*, *80*, 2017, pp.148-174.
[21] Kelly, F.P., "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, Vol. 8, pp. 33–37, January 1997.
[22] Eryilmaz, A., and R. Srikant, "Fair Resource Allocation in Wireless Networks using Queue-length-based Scheduling and Congestion Control," Proceedings of *IEEE Infocom*, 2005.
[23] Corbett-Davies, S., Pierson E., Feller A., Goel, S., and Huq, A., "Algorithmic decision making and the cost of fairness," in *Proceedings of KDD '17*, Halifax, NS, Canada, August 13-17, 2017. doi: 10.1145/3097983.3098095
[24] Young, H.P., *Equity in Theory and Practice*, Princeton University Press, 1994.
[25] Antoine, J., Dubot, T., and Bedouet, J., "Towards a 4D traffic management of small UAS operating at very low level," *ICAS, 30th Congress of the International Council of the Aeronautical Sciences*, 2016.
[26] Jonas, L., Palmerius, K. L., and Josefsson, B., "Urban Air Traffic Management (UTM) Implementation In Cities-Sampled Side-Effects," *2018 IEEE/AIAA 37th Digital Avionics Systems Conference (DASC),* IEEE, 2018.
[27] Li, S., Egorov, M., & Kochenderfer, M. J., "Optimizing Collision Avoidance in Dense Airspace using Deep Reinforcement Learning," *Thirteenth USA/Europe Air Traffic Management R&D Seminar*, Vienna, Austria, 2019.
[28] Sunil, E., Hoekstra, J., Ellerbroek, J., Bussink, F., Nieuwenhuisen, D., Vidosavljevic, A., and Kern, S., "Metropolis: Relating airspace structure and capacity for extreme traffic densities," In *ATM seminar 2015, 11th USA/EUROPE Air Traffic Management R&D Seminar*, 2015.
[29] Ong, H. Y., & Kochenderfer, M. J., "Markov decision process-based distributed conflict resolution for drone air traffic management," *Journal of Guidance, Control, and Dynamics*, 69-80, 2016.
[30] Sunil, E., *et al.*, "Metropolis: Relating airspace structure and capacity for extreme traffic densities," *ATM seminar 2015, 11th USA/EUROPE Air Traffic Management R&D Seminar*, 2015.
[31] Wollkind, S., Valasek J., and Ioerger, T., "Automated conflict resolution for air traffic management using cooperative multiagent negotiation," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, 2004.
[32] Jennings, N.R., *et al.*, "Automated negotiation: prospects, methods and challenges," *International Journal of Group Decision and Negotiation*, 199-215, 2001.
[33] del Pozo de Poza, I., "Assessment of fairness and equity in trajectory based air traffic management," Ph.D. Dissertation, University of Glasgow, 2012, URL: http://theses.gla.ac.uk/3237/ [retrieved 27 November 2019].
[34] del Pozo de Poza, I., Vilaplana Ruiz, M.A., and Goodchild, C., "Assessing Fairness and Equity in Trajectory Based Operations," *9th AIAA Aviation, Technology, Integration, and Operations Conference*, Hilton Head, SC, September 21-23, 2009.
[35] Egorov, M., Kuroda, V., Sachs, P., "Encounter Aware Flight Planning in the Unmanned Airspace," *Integrated Communications, Navigation and Surveillance Conference*, Herndon, VA, 2019.